Comparative Analysis of Static and Dynamic Probabilistic Risk Assessment

Christopher J. Mattenberger, Science and Technology Corporation
Donovan L. Mathias, NASA Ames Research Center
Susie Go, NASA Ames Research Center

Key Words: Dynamic PRA, Risk-Informed Design, Space Exploration, Crewed Spacecraft

SUMMARY & CONCLUSIONS

This study examines three different methodologies for producing loss-of-mission (LOM) and loss-of-crew (LOC) risks estimates for probabilistic risk assessments (PRA) of crewed spacecraft. The three bottom-up, component-based PRA approaches examined are a traditional static fault tree, a dynamic Monte Carlo simulation, and a fault tree hybrid that incorporates some dynamic elements. These approaches were used to model the reaction control system thruster pod of a generic crewed spacecraft and mission, and a comparative analysis of the methods is presented.

The methodologies are assessed in terms of the process of modeling a system, the actionable information produced for the design team, and the overall fidelity of the quantitative risk evaluation generated. The system modeling process is compared in terms of the effort required to generate the initial model, update the model in response to design changes, and support mass-versus-risk trade studies. The results are compared by examining the top-level LOM/LOC estimates and the relative risk driver rankings at the failure mode level. The fidelity of each modeling methodology is discussed in terms of its capability to handle real-world system dynamics such as cold-sparing, changes in mission operations due to loss of redundancy, and common cause failure modes.

The paper also discusses the applicability of each methodology to different phases of system development and shows that a single methodology may not be suitable for all of the many purposes of a spacecraft PRA. The fault tree hybrid approach is shown to be best suited to the needs of early assessments during conceptual design phases. As the design begins to mature, the level of detail represented in the risk model must go beyond redundancy and nominal mission operations to include dynamic, time- and state-dependent system responses as well as diverse system capabilities. This is best accomplished using the dynamic simulation approach, since these phenomena are not easily captured by static methods. Ultimately, once the design has been finalized and the goal of the PRA is to provide design validation and requirement verification, more traditional, static fault tree approaches may become as appropriate as the simulation method.

1 INTRODUCTION

Implementation of risk-informed design allows a design team to thoroughly explore the risks of a system while iterating the operations concept, design, and requirements until the system meets mission objectives and constraints [1]. To arrive at a space system design that is likely to meet all constraints placed upon mass, cost, performance, and risk, the system requirements must be understood and traded against each other as early as the conceptual design phase [2]. Depending on the project phase and the goals of the risk analysis, various PRA methodologies could be used to produce quantitative risk estimates supporting such a process.

In order to better understand the applicability, advantages, and limitations of various PRA methodologies, a comparative analysis of three bottom-up, component-based PRA approaches was performed. The three approaches examined are a traditional static fault tree, a dynamic Monte Carlo simulation, and a fault tree hybrid that incorporates some dynamic elements. Each approach was used to assess a generic reaction control system (RCS) thruster pod and mission [3]. The methods are assessed in terms of the process of modeling a system, the actionable information produced for the design team, and the overall fidelity of the quantitative risk evaluation generated. The paper also discusses the applicability of each methodology to the different phases of system development.

2 REACTION CONTROL SYSTEM DESCRIPTION

The nominal mission under consideration is that of a crewed spacecraft visiting the International Space Station (ISS). The spacecraft is launched into orbit and then must use its onboard propulsion and RCS to rendezvous and dock with ISS 24 hours after launch. Once docked, the spacecraft remains on orbit for 210 days while the RCS is relatively quiescent. Once the spacecraft has completed its stay at ISS, or in the event of an abort from orbit, the spacecraft must once again use its propulsion and RCS to perform de-orbit, entry, descent, and landing operations to return the crew safely within 4 hours.

The RCS thruster pod considered consists of two groups of three thrusters. Loss of any two thrusters in the same group triggers an abort from orbit and ends the nominal mission, thus producing a loss of mission (LOM), and the loss of an entire group triggers a loss of crew (LOC). A simplified schematic of this system is shown in Figure 1, with thrusters represented as blue triangles and isolation valves represented as blue boxes. The nominal operation of the system calls for a "stand-by" thruster-firing protocol, with Thruster A to be fired until it experiences a failure, then Thruster B is fired until failure, and finally Thruster C would then be used to return the crew.

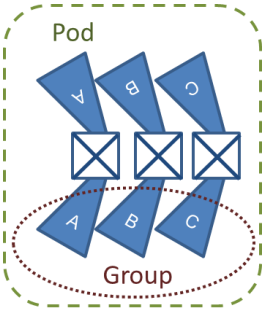

Figure 1 – RCS Thruster Pod Configuration

Each thruster consists of a fuel valve, an oxidizer valve, and an exciter. The valves failure modes are failure to open on demand, failure to close on demand, operational failure while firing, and leakage over time. The exciter can fail off when powered on. The failure rate data is summarized in Table 1. These failure rates are from tables provided in the Institute of Electrical and Electronics Engineers reliability data book [4].

Failure Mode	Failure Rate / Probability
Valve fails to open	2.05e-6 / demand
Valve fails to close	1.51e-6 / demand
Valve fails operationally	9.00e-7 / hour
Valve fails leaky	5.00e-8 / hour
Exciter fails off	1.69e-5 / hour

Table 1 – Failure Rate Data

Each fuel and oxidizer valve is backed up by an isolation valve, which is shared between two thrusters in different groups. The isolation valve is nominally open and only closes if one of two downstream valves has failed open or leaks. If the isolation valve fails to close or leaks, then a LOC is assumed to occur immediately. If the isolation valve successfully closes, then both downstream thrusters are deselected for the rest of the mission and the isolation valve must not leak in order to avoid either LOC while crewed or loss of vehicle (LOV) while docked to ISS.

Table 2 summarizes the risk exposure times and demands by mission phase for each thruster group in the pod. It is important to note that only the currently selected thruster is demanded to fire, while all other thrusters only accrue risk of a valve leakage failure. This leads to uncertainty about how many demands and how much firing time will be accrued by each individual thruster in this cold-spare configuration during an actual mission.

Failure Mode	Pre-Docking	Docked	Post-Undock
Valve fails to open	2,000 demands	N/A	1,000 demands
Valve fails to close	2,000 demands	N/A	1,000 demands
Valve fails operationally	2 hrs	N/A	1 hr
Valve fails leaky	24 hrs	5,040 hrs	4 hrs
Exciter fails off	24 hrs	N/A	4 hrs

Table 2 – Risk Exposure by Mission Phase

Table 3 gives the common cause factor (CCF) values that were used for the RCS thrusters. These CCFs are based upon the Global Alpha Modeling Tool (GAMUT) [5] developed at NASA Johnson Space Center. The values assume that the thrusters use demand-type components that undergo a staggered testing scheme in which they are periodically inspected for indications of incipient failure modes.

Common Cause Group Size	Common Cause Factor
CCF of 2 out of 3	0.04830
CCF of 3 out of 3	0.00517

Table 3 – Common Cause Failure Conditional Probabilities

3 *METHODOLOGIES*

Despite the common set of assumptions about concept of operation, risk exposure, and failure rates presented here, each method must make additional assumptions in order to produce a risk estimate. As such, each methodology provides a risk estimate for an approximate problem. The degree of the approximation versus the cost of obtaining the solution, in terms of risk analyst effort and time, is of key interest in determining the value provided to the design team.

3.1 *Static Fault Tree Approach*

This approach utilized SAPHIRE 8 [6], developed at the Idaho National Laboratory, to construct a static fault tree of the RCS thruster system risks. Multiple instances of the fault tree were constructed to capture the various LOM and LOC end states, since a single model cannot capture both. Having multiple models of the same system can prove difficult to manage if the design is rapidly evolving, the turnaround time for performing trade studies is fast, or the inputs are in flux.

The basic events of the fault tree were calculated off-line and loaded into the model. A major assumption that must be made is determining how many demands each thruster must undertake successfully. Conservatively, it could be assumed that each thruster in a group must fire all 3,000 demands of the mission. However, this excessive conservatism produces unrealistically high risk estimates that are not useful. For this assessment it has been optimistically assumed that all three thrusters in a group each fire an equal amount. The true dynamic reallocation of numbers of thruster firings, firing times, and leakage times cannot easily be accounted for in a

fault tree. For example, an isolation valve should only begin to accrue leakage risk after a random thruster valve failure, but because the time of this failure is uncertain, the model must conservatively assume that the isolation valve must not leak for the entire mission duration.

Common cause failure modes are only captured when they would result directly in a LOM or LOC, depending upon the end-state of the model. Thus, the model does not take into account mixed cases of both random and common cause failure modes combining to cause LOC.

The fault tree approach also does not allow for an elegant method of accounting for time-varying abort criteria. In order to capture these time- and state-dependent system functionalities and behaviors, an intractable number of event trees and corresponding fault trees would need to be constructed. This would make the assessment prohibitively costly and unable to keep up with a rapidly evolving conceptual or preliminary design. However, using this type of method with a long development lead-time and conservative assumptions may be appropriate later in the critical design phase when the design has stabilized and the purpose of the assessment is to verify that it meets a risk requirement.

After the model is created, it must be solved using a specific method in the SAPHIRE program. Both the results and the computation time can vary widely, depending on the chosen solver method and the number of cut-sets it produces. The cut-sets capture all of the model's possible failure modes and their calculated probabilities deterministically, yielding an incredible amount of data that must be processed in order to provide actionable information to decision makers.

Overall, this method produces a very precise solution, but to a very approximate problem. It is extremely useful for rigorously capturing all potential failure modes of a static approximation of the system, but suffers from a lack of responsiveness, which can be a detriment in assessing rapidly evolving designs.

3.2 Dynamic Monte Carlo Simulation Approach

This approach utilizes commercially available, Monte Carlo-style simulation software called GoldSim [7]. The approach uses more complex models that seek to include all dynamic interactions and dependencies between all components and failure modes.

In addition to LOM and LOC estimates, these models are also able to produce estimates of LOV or crew-stranding at ISS, scenario-based event timing information, and data on successful missions or degraded vehicle states that do not trigger LOC, LOV, or LOM. These results can provide decision makers with great insight into maintenance concerns or the value of repair capability.

Unlike a traditional fault tree, the dynamic approach is able to handle more complex, and often more representative, graph-like connections and dependencies that occur in many space systems. The Monte Carlo approach inherently allows dynamic reallocation of demands and changes in system topology that may occur after a failure. Common cause failure modes can be gracefully introduced into the model framework, enabling complete simulation of the Multiple Greek Letter (MGL) [8] method, which represents CCF phenomena more accurately than other approaches.

Overall, this method most accurately captures the system's behavior and yields the greatest design insights, but comes at the cost of greatly increased model complexity. This complexity reduces the model's ability to rapidly respond to an evolving design, makes debugging and validation extremely challenging, and increases computational run-times depending on the number of realizations required to achieve the desired level of confidence in the risk estimate. These factors can make the approach too costly to effectively support risk-informed design in early stages of development. However, recent advancements in cloud computing [9] are reducing the time required to produce risk estimates at the desired confidence level and may enable these complex analysis techniques to become advantageous earlier in the design cycle.

3.3 Rapid Fault Tree Hybrid Approach

The hybrid approach uses the Ames Reliability Tool (ART), which is an Excel-based, implicit event-tree/fault-tree generator developed at NASA Ames Research Center based upon previous work [10]. The ART model deterministically produces estimates of LOM and LOC while capturing some the system's dynamic elements that are expected to drive risk. The ART focuses on risk-driving cut-sets, which are expected to be those due to common cause failure modes and combinatorial mixed cases of both random failures and common cause failures within a specific failure mode or component.

The ART model is able to capture dynamic reallocation of demands after a failure by using well-known "cold spare" or stand-by unit redundancy calculations [11]. This method also accounts for the dynamic change in mission duration if an abort is triggered, and accurately reflects the reduction in crew risk in the case of a degraded vehicle state. Additionally, only one model of the system needs to be built, as the ART is able to produce both LOM and LOC estimates with an extremely simple set of input fields. This method utilizes the built-in functionality of the ART to rapidly create and update models, enabling the risk analyst to work in real-time with designers.

The limitations of this method are that the ART is not able to handle all potential redundancy configurations and does not take into account cross-cutting failure modes between different types of components or different failure modes within a set of similar components. As such, this method does not account for cascade failure modes where a thruster failing open in one group propagates to deselect the corresponding thruster in the other group due to activation of the shared isolation valve. Furthermore, the ART model does not capture thruster loss due to combinations of failure modes, i.e., when one thruster fails to open while another thruster fails to close.

Depending on the system's risk-driving failure modes, the overall risk results may or may not be impacted by optimistically omitting these cut-sets since they contain only random failures, which are often lower probability than those

containing common cause failures. Moreover, if the purpose of the risk assessment is to compare two competing designs, then it is conceivable that these failure modes will not be a difference that makes a difference in the design trade study.

This method sacrifices precision in the absolute risk estimate in order to respond more rapidly to the needs of the decision makers. It captures the system's key dynamics to provide accurate relative rankings of the risk drivers. It also allows the risk analyst to quickly produce a range of estimates based upon uncertain input data to determine the sensitivity of the estimate to the lack of design knowledge.

4 RESULTS

LOM and LOC results from the three modeling approaches are presented in Figure 2. As expected, the hybrid model predicts lower risk than the simulation methodology due to its known omission of cross-failure or cross-component failure modes. The hybrid model captures the majority of the LOM and LOC risks, which stem from CCF modes.

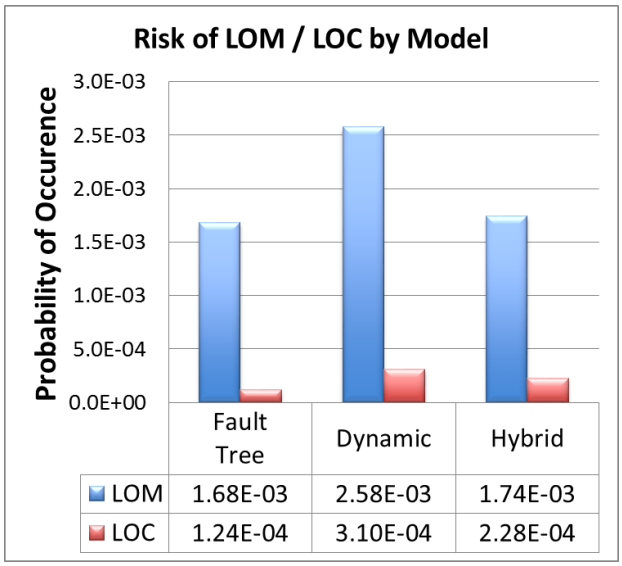

Figure 2 – System-Level LOM/LOC Results by Model

The fault tree results were calculated with both the 'Min-Cut' and 'BDD' solvers built into the SAPHIRE program, which produced numerical results that differed by 2%. The fault tree results are lower than those of the other methods due to the optimistic assumption about the duty cycle for each thruster, which was necessitated by the inability to capture dynamic demand reallocation. Such an approach does not take into account the additional demands that other thrusters must undertake to make-up for those of a failed group member. Conversely, conservatively neglecting to include dynamic abort modes in the fault tree resulted in higher LOC risk.

The dynamic approach results were obtained running 100,000 Monte Carlo simulations over a period of 11 hours on a quad-core Intel i5 processor. Determining the proper number of realizations is important to achieving converged results at the desired level of confidence. Producing high-fidelity results that capture all possible component connectivity and dynamic reallocation of RCS demands would require a prohibitive number of realizations. However, this degree of fidelity is not necessary to obtain a converged estimate at the system level.

All of the methods considered can produce a top-level estimate of LOM and LOC. However, both the hybrid and fault tree approach must make many assumptions to approximate the real-world system. As a result, when compared to the dynamic approach, the hybrid approach underestimates LOM and LOC by 33% and 26%, respectively, and the fault tree underestimates LOM and LOC by 35% and 60%, respectively. Depending on the degree of dynamics and graph-like component interactions, such assumptions could introduce so much uncertainty into the results that they would provide minimal actionable information. If the omissions in modeling fidelity drive the system risk, then the results cannot be trusted on an absolute scale and relative risk results between competing design options cannot be utilized.

A major benefit of the simulation method is that it also records the time at which failure occurs. Such information can be extremely useful if the consequences of failure are time- and state-dependent, such as during an ascent to orbit on a failing launch vehicle, or if increased time on orbit would enable additional scientific research and increased availability of the ISS. In particular, accounting for failure timing allows failures that occur while docked to ISS to be counted as LOV instead of LOC. The dynamic results can provide insight into degraded system states that do not lead to a LOM, but simply to a loss of redundancy and continuation of the nominal mission. This, in turn, can inform expected component failure frequencies to aid in determining repair capabilities and maintenance schedules.

Point estimates of system-level risk can be useful for comparing two different design options or determining if a design meets requirements. However, during the conceptual and preliminary phases of system development, insights into the system's current risk drivers can provide designers with valuable guidance and feedback on how to most effectively increase system reliability and safety. The LOM risk drivers at the thruster failure mode level are provided in Figure 3 for the fault tree method and in Figure 4 for the hybrid and dynamic methods. The fault tree model results for these cases do not immediately yield actionable information to design teams, as the LOM and LOC fault trees respectively produced 900 and 3,956 cut-sets of exact system failure modes, precisely capturing data about which components failed in what mode. For LOM, there are only 24 unique classes of cut-sets when the specificity of which exact component failed is removed. However, it is still difficult to directly apply the results in Figure 4 to provide actionable information to the design team.

Similarly, the results from the dynamic simulation method must also be processed. Figure 3 shows the frequency of failure of each component during a mission with a LOM outcome. Both the overall frequency of failure and frequency of failure leading to a LOM provide actionable information about what failure modes are driving the system risk. However, a drawback of the dynamic method is that the true frequencies of failure modes not observed during any of the simulation realizations remain somewhat uncertain. In this

case, there were no observed failures of the isolation valves, even though this failure mode does show up in both the hybrid and fault tree results. This failure is over-represented in both the fault tree and hybrid models due to their conservative assumption that the diverse leak protection provided by the isolation valves must be reliable for the entire mission, since they are not able to capture this dynamic behavior explicitly.

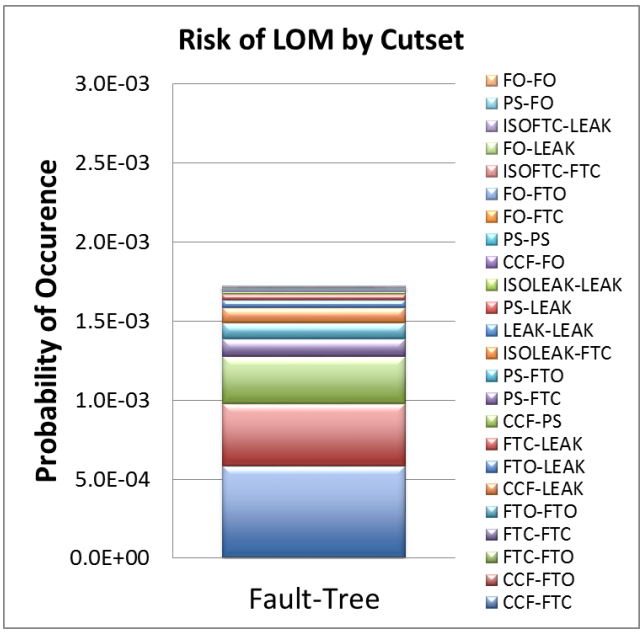

Figure 3 – LOM Results for the Fault Tree Method

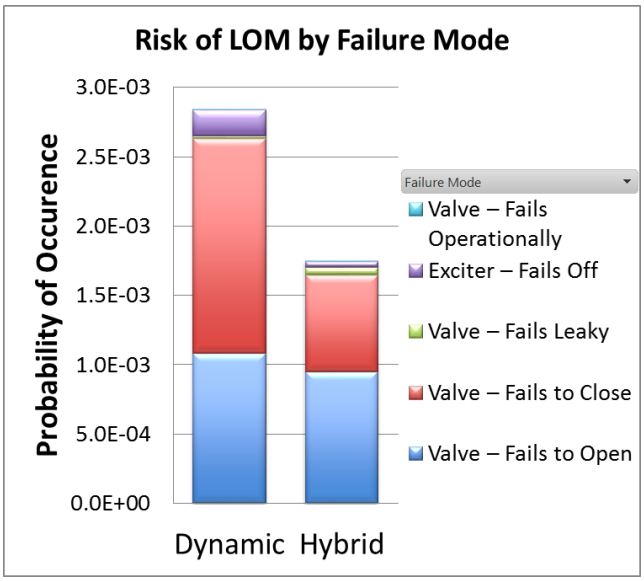

Figure 4 – LOM Results for the Dynamic and Hybrid Methods

The hybrid model immediately provides an ordered list of approximate risk drivers at the failure mode level, but since the approach neglects to model cross-component interactions, it omits an absolute portion of the risk. Since these cross-component cut-sets do not drive the system risk, however, the primary relative risk drivers remain the same.

Providing risk data at the failure-mode level can yield much richer insights into how system safety and reliability can be improved most efficiently. In this case, it is clear that the dominant failure modes are valves failing to open or valves failing to close. The fault tree results do provide the additional, beneficial information that it is common cause failures of these failure modes that drive system risk. Thus, a designer would want to spend precious project resources, such as mass, to protect against these failure modes by backing up these functions redundantly or reducing the susceptibility of these components to common cause failures.

Overall, the results are driven by common cause failures of the fails-to-open and fails-to-close thruster failure modes. As such, it is interesting to analyze the sensitivity of the results to changes in the common cause factors. Figure 5 shows updated model results with the common cause factors reduced by approximately an order of magnitude.

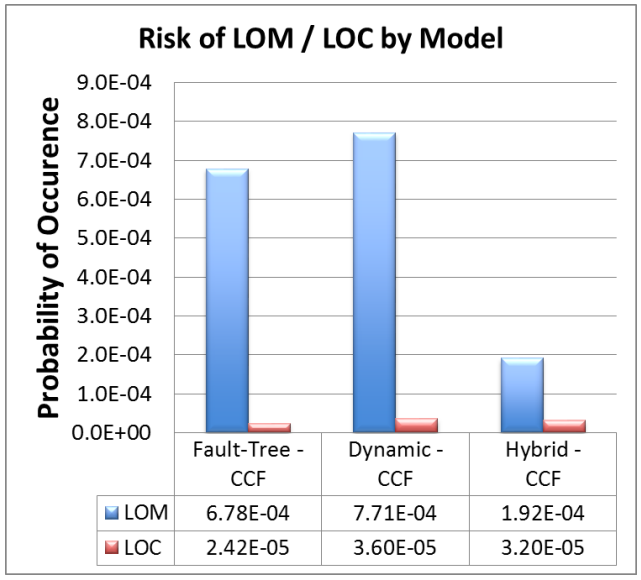

Figure 5 – System-Level LOM/LOC Results by Model with Reduced Common Cause Factors

The hybrid results for LOM now only capture 25% of the risk estimated by the dynamic model since the system risk of LOM is now driven by mixed combinations of random component failures and failure modes rather than by common cause failures. However, the hybrid model does capture 89% of the LOC risk estimated by the dynamic model, as this risk is driven by common cause failure of an entire group as well as by common cause failure of two components and a random failure during the abort.

The fault tree model captures 88% of the LOM risk estimated by the dynamic model, as it does not correctly account for the reallocation of demands to the remaining thrusters after a single random failure. Moreover, the fault tree model only captures 67% of the LOC risk estimated by the dynamic model, since it does not consider the common cause failure of two thrusters triggering an abort and abort modes are not considered in this approach.

5 DISCUSSION

The applicability of each methodology to the different phases of system development can now be discussed in light of the benefits and drawbacks presented here.

The methodology selected during the conceptual design phase needs to respond rapidly to a changing design and provide accurate relative risk drivers given limited design detail. The methodology best suited to providing such insights is the rapid fault tree hybrid approach. Interestingly, many of the limitations associated with the hybrid approach are minimized early in the system development life cycle because the precise design details about cross-strapping and component connectivity are still yet to be determined. Moreover, the purpose of PRA during the conceptual design phase is to guide initial design decisions. Thus, most PRA in this phase will be of a relative nature and a precise, absolute risk estimate is not as important as the comparative differences between multiple, competing design options. Furthermore, at this phase of development, the PRA is often more concerned with reliability potential rather than "as-drawn" reliability.

As the design begins to mature, more precise insights are required to accurately discriminate between similar trade study options and identify the factors that can most efficiently reduce overall risk. Additionally, providing accurate, absolute risk estimates becomes increasingly important to enable comparison of design options in unrelated subsystems. Design trade studies start to become more subtly nuanced and require precise representations of real-world system operation. To accomplish this, the level of design detail represented in the risk model must go beyond redundancy and nominal mission operations to include dynamic, time- and state-dependent system responses as well as diverse system capabilities. The dynamic simulation methodology is best suited to this phase of development, as many risk-driving and risk-differentiating phenomena are not easily captured by static methods.

Ultimately, once the design has been finalized, more traditional, static fault tree approaches may become as appropriate as the simulation method. At this point in the design cycle, the goal of the PRA is often to show that the system meets requirements or to validate the design by exhaustively searching for unintended failure modes or cut-sets that are not intuitively obvious but are easily revealed through a fault tree. In this case, making overly conservative assumptions can be completely valid. Moreover, since the questions being asked of the PRA are much broader and less specific, the PRA does not have to provide decision makers with as much detailed insight in such a rapid fashion.

REFERENCES

1. J. Miller, J. Leggett, and J. Kramer-White, "*Design Development Test and Evaluation Considerations for Safe and Reliable Human Rated Spacecraft Systems*," NASA, 2008, Hampton, VA.
2. J. R. Fragola, "*Supporting Preliminary Design Decision Making with a Risk Data Base*," Proceedings of Probabilistic Safety Assessment and Management conference, 2010, Seattle, WA.
3. S. A. Motiwala, D. L. Mathias, C. J. Mattenberger, "*Conceptual Launch Vehicle and Spacecraft Design for Risk Assessment*", NASA/TM-2014-218366, NASA ARC, Moffett Field, CA, 2014.
4. L. E. Booth, "*IEEE Guide to the collection and presentation of Electrical, Electronic, Sensing Component and Mechanical Equipment Reliability data for nuclear-power generating stations,"* IEEE Std 500-1984, New York, NY, 1983.
5. Reistle, Bruce, Global Alpha Model Uncertainty Tool (GAMUT), July 2011.
6. Idaho National Laboratory (INL©), SAPHIRE®, Version 8, Idaho Falls, Idaho.
7. www.goldsim.com
8. M. Mohammad, "*Risk Analysis in Engineering: Techniques, Tools, and Trends*", Taylor & Francis Group, Boca Raton, FL, 2006.
9. https://aws.amazon.com/what-is-cloud-computing/
10. B. F. Putney, E. Tavernetti, J.R. Fragola, and E. Gold, "*Reliability Tool for a Preliminary Quantified Functional Risk and Hazard Analysis*," Proceedings of the Reliability and Maintainability Symposium, 2009, Fort Worth, TX.
11. D. B. Kececioglu, "*Reliability Engineering Handbook, Volume 2*", DEStech Publications, Lancaster, PA, 2002.

BIOGRAPHIES

Chris Mattenberger
NASA Ames Research Center, Mail Stop 258-6
Moffett Field, CA 94035-1000 USA

email: chris.mattenberger@nasa.gov

Mr. Mattenberger is a Research Scientist/Engineer with the Science and Technology Corporation. He received a B.S. in Aerospace Engineering from the Massachusetts Institute of Technology, and an M.S. in Aeronautics and Astronautics from Stanford University.

Donovan L. Mathias
NASA Ames Research Center, Mail Stop 258-5, Rm. 202
Moffett Field, CA 94035, USA

email: donovan.mathias@nasa.gov

Donovan Mathias is an Aerospace Engineer in the NASA Advanced Supercomputing Division at NASA Ames Research Center. Dr. Mathias earned his B.S. and M.S. in Aeronautical Engineering from the California Polytechnic State University, San Luis Obispo and his Ph.D. in Aeronautics and Astronautics from Stanford University.

Susie Go
NASA Ames Research Center, Mail Stop 258-6
Moffett Field, CA 94035-1000 USA

email: susie.go@nasa.gov

Susie Go is an Aerospace Engineer at NASA Ames Research Center. Dr. Go received her B.A. in the fields of Mathematics and Microbiology/Immunology from the University of California, Berkeley, and her M.A. and Ph.D. degrees in Applied Mathematics from the University of California, Los Angeles.